arXiv:1702.04165v1 [astro-ph.IM] 14 Feb 2017

Crossmatching variable objects with the *Gaia* data

Lorenzo Rimoldini,[1] Krzysztof Nienartowicz,[2] Maria Süveges,[1]
Jonathan Charnas,[1] Leanne P. Guy,[1] Grégory Jevardat de Fombelle,[2]
Berry Holl,[1] Isabelle Lecoeur-Taïbi,[1] Nami Mowlavi,[3]
Diego Ordóñez-Blanco,[1] and Laurent Eyer[3]
[1]*Department of Astronomy, University of Geneva, Chemin d'Ecogia 16, CH-1290 Versoix, Switzerland;* Lorenzo.Rimoldini@unige.ch

[2]*SixSq, Rue du Bois-du-Lan 8, CH-1217 Geneva, Switzerland*

[3]*Department of Astronomy, University of Geneva, Chemin des Maillettes 51, CH-1290 Versoix, Switzerland*

Abstract. Tens of millions of new variable objects are expected to be identified in over a billion time series from the *Gaia* mission. Crossmatching known variable sources with those from *Gaia* is crucial to incorporate current knowledge, understand how these objects appear in the *Gaia* data, train supervised classifiers to recognise known classes, and validate the results of the Variability Processing and Analysis Coordination Unit (CU7) within the *Gaia* Data Analysis and Processing Consortium (DPAC). The method employed by CU7 to crossmatch variables for the first *Gaia* data release includes a binary classifier to take into account positional uncertainties, proper motion, targeted variability signals, and artefacts present in the early calibration of the *Gaia* data. Crossmatching with a classifier makes it possible to automate all those decisions which are typically made during visual inspection. The classifier can be trained with objects characterized by a variety of attributes to ensure similarity in multiple dimensions (astrometry, photometry, time-series features), with no need for a-priori transformations to compare different photometric bands, or of predictive models of the motion of objects to compare positions. Other advantages as well as some disadvantages of the method are discussed. Implementation steps from the training to the assessment of the crossmatch classifier and selection of results are described.

Introduction. The crossmatch of celestial objects makes it possible to combine complementary information from data collected at various epochs, with different observational and instrumental features (such as wavebands, time sampling, duration, sky coverage, photometric and astrometric accuracy), and also to extract new information by leveraging the synergy among data sets. At the same time, some of the differences in instrumentation and data taking, convolved with the properties of the objects to crossmatch (herein named targets), can lead to misses and false detections (Gray et al. 2007), which can become numerous as the number of targets grows. Common causes of crossmatch errors include large positional uncertainties, proper motion, variability, blended objects, spurious sources, detector edges or gaps, contamination, noise, etc. Variable objects can be more challenging to crossmatch than constant sources, but they also provide additional features which can be exploited to aid in the identification of correct matches. For each object, we consider multiple characteristics derived from astrome-

try, photometry, and light curves, in combination with additional information from literature. Machine-learning classifiers are convenient tools to handle multi-dimensional tasks, automate the variety of decisions common in visual inspections, and minimize the occurrence of false positives and negatives. Supervised classifiers have previously been used for crossmatching catalogues with large positional uncertainties (Richards et al. 2012). Inspired by this work, we extended the classifier method to make full use of the time series information and applied it to crossmatch variable sources in the *Gaia* data with a selection of surveys, for use in validation and training of variability types.

Method. The main steps to crossmatch with a classifier are outlined below and followed by a brief summary of the pros and cons of the method.

Neighbours. The first step to crossmatch a set of targets in a data set is to find the corresponding neighbours in another data set within some angular radius from the targets, after making sure that the coordinates of the two data sets are compared in the same reference system, defined with the same equinox, and possibly taking into account the epoch of observation if the displacement by proper motion over time is not negligible (which might imply a search radius much greater than commonly used). We searched for neighbours with efficient PostgreSQL queries making use of the Quad Tree Cube sky indexing scheme (Koposov & Bartunov 2006). Our search radius was limited to 5 arcsec, accounting mostly for positional uncertainties of ground-based surveys, as most targets were located in the Large Magellanic Cloud (and thus with negligible proper motion effects).

Match criteria. Classification attributes are computed to distinguish matches from non-matches with several criteria from astrometry (angular separation), photometry (e.g., mean brightness, colour), time-series parameters (e.g., central moments and other statistics characterizing the variability), which also incorporate results from literature like periodicity (light curves folded by their most significant periods can be compared effectively, e.g. by the phases of brightness extrema or by a reduced point-to-point scatter). Depending on the criterion, it can be useful to include values as computed in each data set as well as from their comparison (differences or ratios). While the classifier should identify correct matches without relying on positions, if proper motion is relevant, its value could be correlated by the classifier with the angular separation from the target and thus reduce the risk of contamination with similar-looking neighbours.

Training objects. The selection of objects for the training set is one of the most critical phases of supervised classification. To ensure reliable results, a special effort is made to: *(i)* provide a good representation of all match and non-match criteria as a function of variability type and data quality (if the classifier gives different weights to classes depending on their relative representation, we suggest to use a similar number of training matches and non-matches); *(ii)* embed all possible reasons which drive visual inspection-based decisions, including as many challenging cases as possible; *(iii)* verify that the misclassification level is low and that the objects among false positives and negatives correspond to acceptable mistakes, or improve the definition of misclassified objects in the training set and iterate until the above-mentioned conditions are met. Occasionally, additional dedicated classifiers might be needed to deal with especially difficult cases (e.g., to recover matches from objects initially classified as non-matches).

Optimisation. For robust results and to avoid model overfitting, the classifier is optimised by its internal parameters (depending on the method) and by selecting an appropriate subset of the most useful classification attributes (e.g., by forward selection or

backward elimination, see Guyon & Elisseeff 2003). Misclassified training-set objects from the optimised model are then assessed as in item *(iii)* of the training-set selection.

Classification. Finally, the classifier model is applied to the objects to crossmatch. In the current version, we assume that only one match is associated with each target and vice-versa. When the *Gaia* observations split sources which are unresolved or blended in other surveys, there is still some chance that the variable object is correctly identified by its variability pattern (unless the system includes multiple variable sources). Crossmatch results might include multiple match candidates per target and matches associated with multiple targets. We decided to select first the highest probability match for each target. If among the selected matches more than one target is associated with the same match, different options are possible: retain the safest matches (keeping the one with the highest probability and then iterating on the remaining targets for the next highest match probability until there are only single targets per match) or aim at crossmatch completeness (including lower match probabilities but for more targets). For the *Gaia* data, we chose to base our selection on the reliability of the crossmatch (based on the highest probability), as presented in Fig. 1.

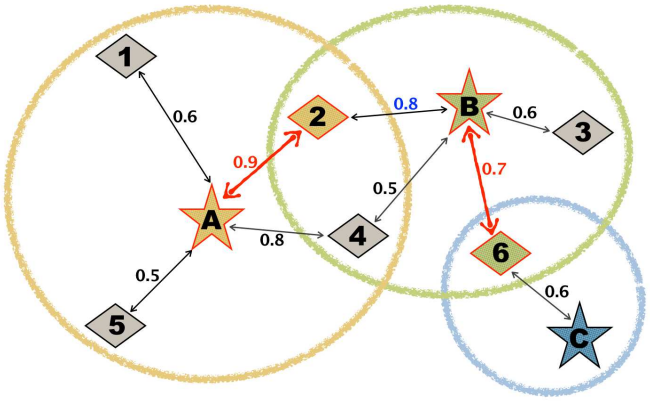

Figure 1. Probability-based selection of the best-match objects. Targets to crossmatch are denoted by letters (A, B, C) and their match candidates by integer numbers (from 1 to 6), while the respective match probabilities are indicated by the values next to the arrows connecting each object to one or more targets. In the cases of multiple matches per target, we select the match-target pairs with the highest probabilites. When the same match is associated with two targets, it is assigned to the target with the highest probability, leaving the other target to another match (with the second best probability), if any. In the case depicted here, target A is associated with object 2, target B with object 6, and target C remains unmatched.

Assessment. Classification results are assessed by inspecting low-probability non-matches and matches, the farthest matches, the nearest non-matches, and other potential border-line cases. While misclassifications are almost inevitable, the cases which cannot be missed are included in the training set (possibly with additional similar objects) and the steps from the classifier optimisation are iterated until misclassifications are acceptable. Further diagnostics of the global results, such as the distribution of matches in magnitude- and colour-difference space, can help highlight issues and direct corrective actions (like the selection of new training set objects and/or attributes).

Pros and cons. In summary, the method of crossmatching with a classifier has several advantages with respect to traditional position-based techniques: *(i)* the ability to characterize objects by a variety of features to better differentiate the match vs. non-match classes and automatically minimize the error rate; *(ii)* robustness of results as the classifier adapts to the data and discovers intrinsic relations: imperfect calibrations do not prevent optimal results, biases caused by artefacts are accounted for (as trained), measurements in different photometric bands can be compared directly without a-priori transformations (as long as the quantities which define them, such as brightness and colour, are included as attributes); *(iii)* better performance than a single multi-dimensional metric, as it does not depend on the accuracy of the components or theoretical expectations; *(iv)* selectivity based on variability: if the variability signals of a target and a match candidate are different, the classifier can be taught to consider the pair as a match (e.g., if the signal is only partially sampled) or non-match (e.g., if there is no interest in an eclipsing binary with no measurement in the eclipses); *(v)* independence from astrometric details: matches with low positional accuracy or significant proper motion can be identified without knowledge of positional uncertainties or predictive models of their positions (as long as they are within the neighbour search radius); *(vi)* the classifier returns a reliability score in the form of an estimate of the probability of matches, which can also be used to set different thresholds depending on the purpose (e.g., a higher threshold for training variability types and a lower one for completeness analyses). On the other hand, the main disadvantage of the supervised classifier method is that it depends on the training set (by definition) and it takes time to select training-set objects properly. As every survey is unique, new classifiers must be trained to crossmatch with different data sets. Considering the time to visually inspect hundreds of sources for a good training set, the visual confirmation of the best match among the neighbours can be more efficient when the number of crossmatch targets is less than about a thousand.

Results. The method described herein was applied by means of Random Forest classifiers (Breiman 2001) to crossmatch known variable objects with the *Gaia* data. Full details of the crossmatch results, crossmatched catalogues, number of matches per catalogue and their sky coverage are presented in Eyer et al. (2017). Crossmatch targets covered primarily the region near the LMC, mostly from the OGLE-IV (Soszyński et al. 2012, 2015a,b) and the EROS-II (Kim et al. 2014; Tisserand et al. 2007) surveys.

References

Breiman, L. 2001, Machine Learning, 45, 5
Eyer, L., et al. 2017, (submitted). 1702.03295
Gray, J., Szalay, A., Budavari, T., Lupton, R., Nieto-Santisteban, M., & Thakar, A. 2007, eprint. cs/0701172
Guyon, I., & Elisseeff, A. 2003, J. Machine Learning Res., 3, 1157
Kim, D.-W., Protopapas, P., Bailer-Jones, C. A. L., Byun, Y.-I., Chang, S.-W., Marquette, J.-B., & Shin, M.-S. 2014, A&A, 566, A43. 1403.6131
Koposov, S., & Bartunov, O. 2006, in ADASS XV, edited by C. Gabriel, C. Arviset, D. Ponz, & S. Enrique, vol. 351 of Astronomical Society of the Pacific Conference Series, 735
Richards, J. W., Starr, D. L., Miller, A. A., Bloom, J. S., Butler, N. R., Brink, H., & Crellin-Quick, A. 2012, Astrophys. J. Suppl. Series, 203, 32. 1204.4180
Soszyński, I., et al. 2012, Acta Astron., 62, 219. 1210.1219
— 2015a, Acta Astron., 65, 233. 1508.00907
— 2015b, Acta Astron., 65, 297. 1601.01318
Tisserand, P., et al. 2007, A&A, 469, 387. astro-ph/0607207